Network Working Group
Internet Draft
Category: Informational
Expires: March 2012

Raymond Key (editor), Huawei
Simon Delord, Alcatel-Lucent
Frederic Jounay, France Telecom
Lu Huang, China Mobile
Zhihua Liu, China Telecom
Manuel Paul, Deutsche Telekom
Ruediger Kunze, Deutsche Telekom
Nick Del Regno, Verizon
Joshua Rogers, Time Warner Cable

September 29, 2011

Requirements for MEF E-Tree Support in VPLS
draft-key-l2vpn-vpls-etree-reqt-04

Status of this Memo

Abstract

This document provides functional requirements for Metro Ethernet Forum (MEF) Ethernet Tree (E-Tree) support in Virtual Private LAN Service (VPLS). It is intended that potential solutions will use these requirements as guidelines.

Table of Contents

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

1. Introduction

This document provides functional requirements for Metro Ethernet Forum (MEF) Ethernet Tree (E-Tree) support in Virtual Private LAN Service (VPLS). It is intended that potential solutions will use these requirements as guidelines.

Considerable number of service providers have adopted VPLS to provide MEF Ethernet LAN (E-LAN) services to customers. Service Providers currently need a simple and effective solution to emulate E-Tree services in addition to E-LAN services on their MPLS networks.

2. Virtual Private LAN Service

VPLS is a L2VPN service that provides multipoint-to-multipoint connectivity for Ethernet across an IP or MPLS-enabled IP Packet Switched Network. VPLS emulates the Ethernet VLAN functionality of traditional Ethernet network.

VPLS is a current IETF standard, please refer to [RFC4761] [RFC4762].

Data frame is Ethernet frame.

Data forwarding is MAC-based forwarding, which includes MAC address learning and aging.

3. MEF Multipoint Ethernet Services

MEF has defined two multipoint Ethernet Service types:
- E-LAN (Ethernet LAN), multipoint-to-multipoint service
- E-Tree (Ethernet Tree), rooted-multipoint service

For full specification, please refer to [MEF6.1] [MEF10.2].

3.1. Similarity between E-LAN and E-Tree

Data frame is Ethernet frame.

Data forwarding can be MAC-based forwarding or something else, to be specified by service provider as service frame delivery attributes in the particular service definition.

A generic E-LAN/E-Tree service is always bidirectional in the sense that ingress frames can originate at any endpoint in the service.

3.2. Difference between E-LAN and E-Tree

Within the context of a multipoint Ethernet service, each endpoint is designated as either a Root or a Leaf. A Root can communicate with

all other endpoints in the same multipoint Ethernet service, however a Leaf can only communicate with Roots but not Leafs.

The only difference between E-LAN and E-Tree is:

- E-LAN has Root endpoints only, which implies there is no communication restriction between endpoints
- E-Tree has both Root and Leaf endpoints, which implies there is a need to enforce communication restriction between Leaf endpoints

3.3. E-Tree Use Cases

Table 1 below presents some major E-Tree use cases.

	Use Case	Root	Leaf
1	Hub & Spoke VPN	Hub Site	Spoke Site
2	Wholesale Access	Customer's Interconnect	Customer's Subscriber
3	Mobile Backhaul	RAN NC	RAN BS
4	IEEE 1588 PTPv2 Clock Synchronisation	PTP Server	PTP Client
5	Internet Access	BNG Router	Subscriber
6	Broadcast Video (unidirectional only)	Video Source	Subscriber
7	Broadcast/Multicast Video plus Control Channel	Video Source	Subscriber
8	Device Management	Management System	Managed Device

Table 1: E-Tree Use Cases

Common to all use cases, direct layer 2 Leaf-to-Leaf communication is not required. For Mobile backhaul, this may not be valid for LTE X2 interfaces in the future.

If direct layer 2 Leaf-to-Leaf communication is not allowed due to security concern, then E-Tree should be used to prohibit communication between Leaf endpoints, otherwise E-LAN is also a feasible option.

3.4. Generic E-Tree Service

A generic E-Tree service supports multiple Root endpoints. The need

for multiple Root endpoints is usually driven by redundancy requirement. Whether a particular E-Tree service needs to support single or multiple Roots depends on the target application.

A generic E-Tree service supports all the following traffic flows:
- Ethernet Unicast from Root to Leaf
- Ethernet Unicast from Leaf to Root
- Ethernet Unicast from Root to Root
- Ethernet Broadcast/Multicast from Root to Roots & Leafs
- Ethernet Broadcast/Multicast from Leaf to Roots

A particular E-Tree service may need to support all the above or only a subset depending on the target application.

4. Problem Statement

4.1. Motivation

VPLS can be used to emulate MEF E-LAN service over MPLS network provided that the E-LAN service uses MAC-based forwarding as service frame delivery attributes.

Considerable number of service providers have adopted VPLS to provide MEF E-LAN services to customers. Service Providers currently need a simple and effective solution to emulate E-Tree services in addition to E-LAN services on their MPLS networks.

4.2. Leaf-to-Leaf Communication Restriction

Current standard VPLS treats all ACs equal (i.e. not classified into Root or Leaf) and provides any-to-any connectivity among all ACs. The current standard VPLS does not include any mechanism of communication restriction between specific ACs, therefore is insufficient for emulating generic E-Tree service over MPLS network.

A problem occurs when there are two or more PEs with both Root AC and Leaf AC.

Let's look at the scenario illustrated in Figure 1 below. VPLS is used to emulate an E-Tree service over a MPLS network.

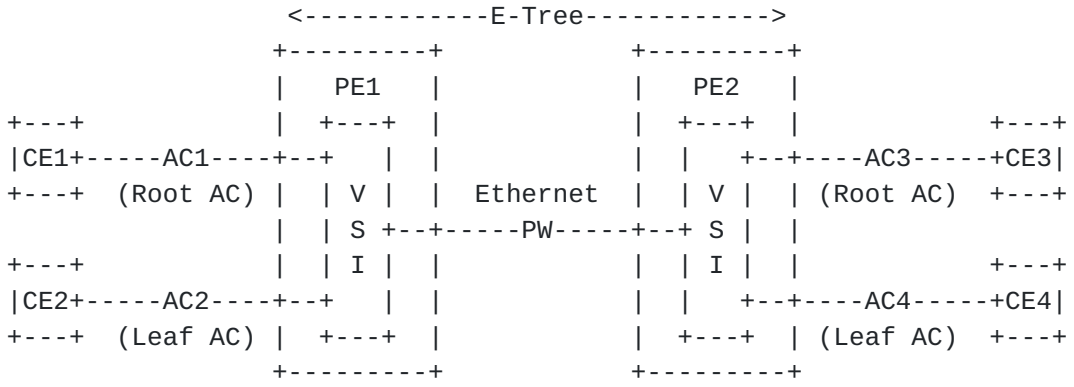

Figure 1: Problem Scenario for Leaf-to-Leaf Communication Restriction

When PE2 receives a frame from PE1 via the Ethernet PW,
- PE2 does not know which AC on PE1 is the ingress AC
- PE2 does not know whether the ingress AC is a Leaf AC or not
- PE2 does not have sufficient information to enforce the Leaf-to-Leaf communication restriction

Examples:
- CE2 sends a Broadcast/Multicast frame to PE1 via AC2
- CE2 sends a Unicast frame to PE1 via AC2, destination address in Ethernet header equal to CE4's MAC address

Note: Figure 1 is a hypothetical case solely for explaining the problem, and not meant to represent a typical E-Tree service.

There are some possible ways to get around this problem that do not require extension to the current standard VPLS but they all come with significant design complexity or deployment constraints, please refer to [Draft ETree Frwk] Appendix A.

5. Requirements

5.1. Functional Requirements

A solution MUST prohibit communication between any two Leaf ACs in a VPLS instance.

A solution MUST allow multiple Root ACs in a VPLS instance.

A solution MUST allow Root AC and Leaf AC of a VPLS instance co-exist on any PE.

5.2. Applicability

There are two distinct VPLS standards, performing similar functions in different manners.

- [RFC4761], commonly known as BGP-VPLS

- [RFC4762], commonly known as LDP-VPLS

A solution MUST identify which VPLS standards the solution is applicable to, [RFC4761] or [RFC4762] or both.

Service providers may use single or multiple technologies to deliver an end-to-end E-Tree service.

- Case 1: Single technology "VPLS Only"

- Case 2: Multiple technologies "VPLS + Others"

- e.g. VPLS + Ethernet network, VPLS + OTN

- Case 3: Single/multiple technologies "No VPLS"
  - e.g. Ethernet network, Ethernet network + OTN
  - out of scope for this document

A solution MUST identify which of the above cases the solution is applicable to. For Case 2, further details may be required to specify the applicable deployment scenarios.

5.3. Backward Compatibility

A solution SHOULD minimise the impact on existing VPLS solution, especially for the MEF E-LAN services already in operation.

A solution SHOULD be backward compatible with the existing VPLS solution. It SHOULD allow a case where a common VPLS instance is composed of both PEs supporting the solution and PEs not supporting it, and the Leaf-to-Leaf communication restriction is enforced within the scope of the compliant PEs.

6. Security Considerations

This will be added in later version of this document.

7. IANA Considerations

This will be added in later version of this document.

8. Acknowledgements

The authors would like to thank Lizhong Jin, Lucy Yong, Yuji Kamite and Wim Henderickx for their valuable input and support.

9. References

9.1. Normative References

[MEF6.1] Metro Ethernet Forum, Ethernet Services Definitions - Phase 2, April 2008

[MEF10.2] Metro Ethernet Forum, Ethernet Services Attributes Phase 2, October 2009

[RFC2119] Bradner, S., Key words for use in RFCs to Indicate Requirement Levels, BCP 14, RFC 2119, March 1997

[RFC4761] Kompella & Rekhter, Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling, January 2007

[RFC4762] Lasserre & Kompella, Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling, January 2007

9.2. Informative References

[Draft ETree Frwk] Key, et al., A Framework for E-Tree Service over MPLS Network, draft-key-l2vpn-etree-frwk-05.txt (work in progress), April 2011

[Draft VPMS Frmwk] Kamite, et al., Framework and Requirements for Virtual Private Multicast Service (VPMS), draft-ietf-l2vpn-vpms-frmwk-requirements-04.txt (work in progress), July 2011

Appendix A. Frequently Asked Questions

A.1. Are E-Tree requirements addressed in the VPMS requirement draft?

VPMS is Virtual Private Multicast Service. VPMS requirement draft refers to [Draft VPMS Frmwk].

The focus of VPMS is to provide point-to-multipoint connectivity.

VPMS provides single coverage of receiver membership (i.e. there is no distinct differentiation for multiple multicast groups). A VPMS service supports single or multiple Root ACs. All traffic from a Root AC will be forwarded to all Leaf ACs (i.e. P2MP, from Root to all Leafs). Destination address in Ethernet frame is not used in data forwarding. As an optional capability, a VPMS service may support reverse traffic from a Leaf AC to a Root AC (i.e. P2P, from Leaf to Root).

In contrast, the focus of MEF E-Tree is that a Leaf can only communicate with Roots but not Leafs.

A generic MEF E-Tree service supports multiple Root endpoints. Whether a particular E-Tree service needs to support single or multiple Root endpoints depends on the target application.

A generic MEF E-Tree service supports all the following traffic flows:

- Ethernet Unicast bidirectional Root to/from Root
- Ethernet Unicast bidirectional Root to/from Leaf
- Ethernet Broadcast/Multicast unidirectional Root to all Roots & Leafs
- Ethernet Broadcast/Multicast unidirectional Leaf to all Roots.

A particular E-Tree service may need to support all the above or only a subset depending on the target application.

IETF's VPMS definition and MEF's E-Tree definition are significantly different.

Only for special case E-Tree service where

- No Unicast traffic from Root destined for a specific Leaf (or there is no concern if such Unicast traffic are forwarded to all Leafs)
- No traffic between Roots

VPMS will be able to meet the requirement. An example is E-Tree service for content delivery application.

For generic E-Tree service, VPMS will not be able to meet the requirements.

A.2. Are there any potential deployment scenarios for a "VPLS Only" solution?

This refers to Section 5.2. Applicability, Case 1: Single technology "VPLS Only".

Yes, there are potential deployment scenarios for a "VPLS Only" solution, some examples below.

Example 1 -

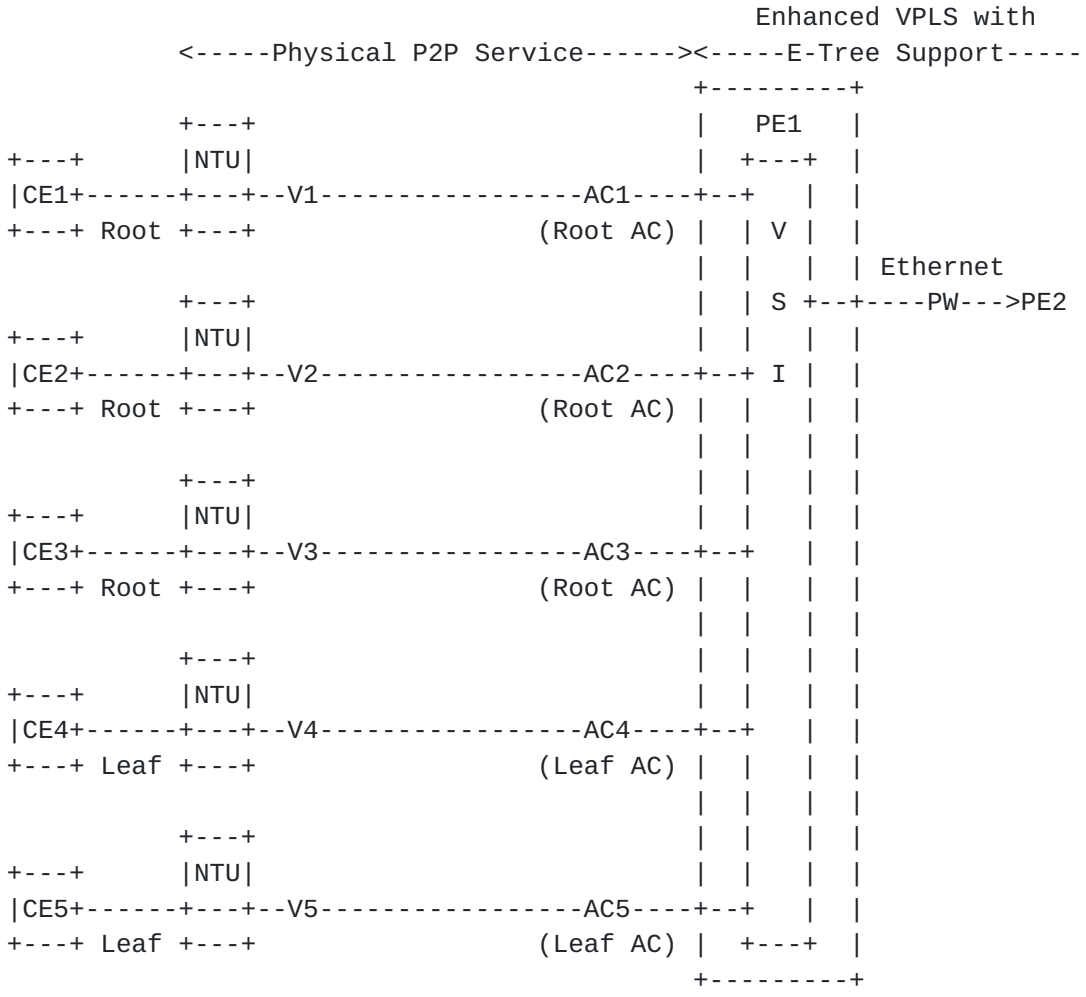

Example 2 -

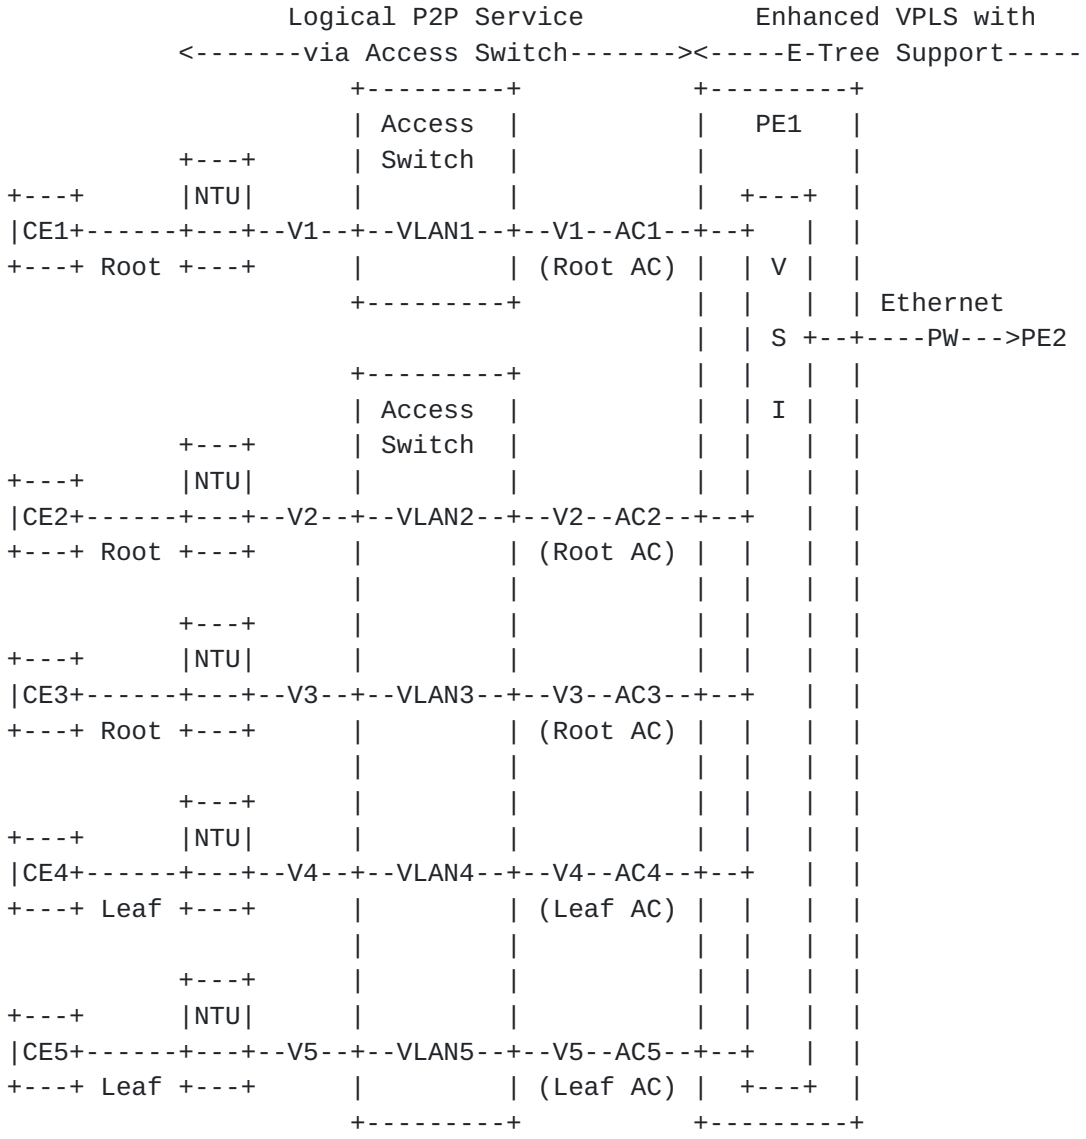
Logical P2P Service via Access Switch
Enhanced VPLS with E-Tree Support
Access Switch
PE1
NTU
CE1
Root
V1
VLAN1
V1--AC1
(Root AC)
V
S
I
Ethernet
PW--->PE2
Access Switch
NTU
CE2
Root
V2
VLAN2
V2--AC2
(Root AC)
NTU
CE3
Root
V3
VLAN3
V3--AC3
(Root AC)
NTU
CE4
Leaf
V4
VLAN4
V4--AC4
(Leaf AC)
NTU
CE5
Leaf
V5
VLAN5
V5--AC5
(Leaf AC)

Example 3 -

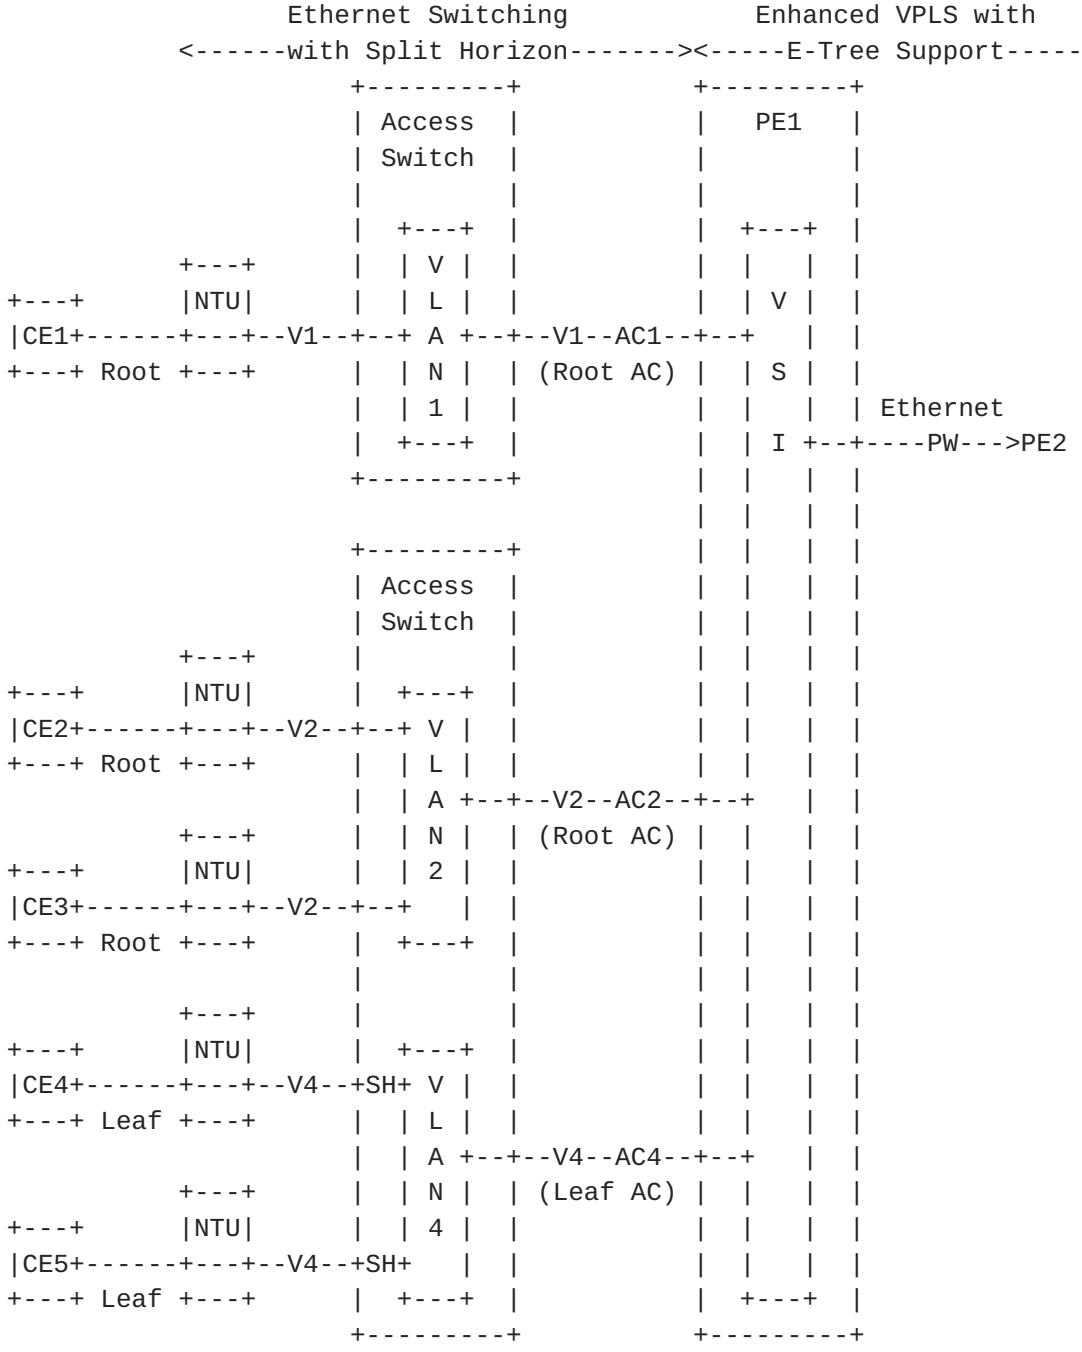

Note:
- Group Roots and Leafs into two separate VLANs on Access Switch
- SH means member of split horizon group on Access Switch

Authors' Addresses

Raymond Key (editor)
Huawei
Email: raymond.key@ieee.org

Simon Delord
Alcatel-Lucent
Email: simon.delord@gmail.com

Frederic Jounay
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex, France
Email: frederic.jounay@orange.com

Lu Huang
China Mobile
Unit 2, 28 Xuanwumenxi Ave, Xuanwu District
Beijing 100053, China
Email: huanglu@chinamobile.com

Zhihua Liu
China Telecom
109 Zhongshan Ave., Guangzhou
510630, China
Email: zhliu@gsta.com

Manuel Paul
Deutsche Telekom
Winterfeldtstr. 21-27
10781 Berlin, Germany
Email: manuel.paul@telekom.de

Ruediger Kunze
Deutsche Telekom
Winterfeldtstr. 21-27
10781 Berlin, Germany
Email: ruediger.kunze@telekom.de

Nick Del Regno
Verizon
400 International Pkwy
Richardson, TX 75081, USA
Email: nick.delregno@verizon.com

Joshua Rogers
Time Warner Cable
11921 N. MoPac Expwy

Austin, TX 78759, USA
Email: josh.rogers@twcable.com

Copyright Notice